

Improved Results for Route Planning in Stochastic Transportation Networks

Justin Boyan
NASA Ames Research Center
Moffett Field, CA 94035
jboyan@arc.nasa.gov

Michael Mitzenmacher*
Harvard University
33 Oxford St.
Cambridge, MA 02138
michaelm@eecs.harvard.edu

Abstract

In the bus network problem, the goal is to generate a plan for getting from point X to point Y within a city using buses in the smallest expected time. Because bus arrival times are not determined by a fixed schedule but instead may be random, the problem requires more than standard shortest path techniques. In recent work, Datar and Ranade provide algorithms in the case where bus arrivals are assumed to be independent and exponentially distributed.

We offer solutions to two important generalizations of the problem, answering open questions posed by Datar and Ranade. First, we provide a polynomial time algorithm for a much wider class of arrival distributions, namely those with *increasing failure rate*. This class includes not only exponential distributions but also uniform, normal, and gamma distributions. Second, in the case where bus arrival times are independent and geometric discrete random variables, we provide an algorithm for transportation networks of buses and trains, where trains run according to a fixed schedule.

*Supported in part by a grant from the Alfred P. Sloan Foundation and an equipment grant from Compaq Computer Corporation.

1 Introduction

Imagine trying to travel across a city by bus, with the goal of minimizing the expected total travel time. There may be several different possible routes, with some requiring changing buses. If buses followed a fixed schedule, then standard shortest-path techniques would be sufficient to find the best travel plan. However, bus arrivals rarely follow a fixed schedule (even when they are supposed to). Bus arrivals are more naturally modeled as a random process, in which case a natural goal is to develop a plan that minimizes the total expected travel time. Although this *bus network problem* appears specific, it is representative of a wide class of scheduling problems where an appropriate plan must be developed with incomplete information modeled probabilistically.

The bus network problem was recently examined by Datar and Ranade, in the case where arrival distributions are independent Poisson processes, i.e. the interarrival times are exponentially distributed, with the mean of each distribution fixed for all time. Their results are based on the key insight that in this case, the optimal plan is composed of statements of the following form: “When at station i , wait for one of buses $X_{i1}, X_{i2}, \dots, X_{ik_i}$; take the first of these buses that arrives.” Moreover, they show that because optimal plans have such a simple form, they can be calculated in polynomial time using a dynamic programming algorithm.

On reflection it is clear that the simple form of the optimal plan is highly dependent on the assumption of independent Poisson arrival processes with fixed means. (We will see examples below.) This assumption is problematic: indeed, the authors admit, “Perhaps the most unfounded assumption in our model is that of Poisson arrivals of the buses.”

As our first result, we show that an optimal plan has only a slightly more complex form when the arrival distributions for buses are assumed to be independent and have *increasing failure rate*. Intuitively, the waiting time for a bus has increasing failure rate if the longer you wait, the more likely the bus is about to arrive. Many natural models—including uniform, normal, and gamma distributions—have increasing failure rate, so our result may be much more appropriate for real-world data. We describe how the optimal schedule in this case can be determined in polynomial time, assuming that we can compute with the relevant probability distributions in an effective manner. Finally, we demonstrate that assuming a slightly weaker property than increasing failure rate for the bus arrival distributions is insufficient for our results.

As our second result, we partially answer another open question posed by Datar and Ranade: how can we handle both buses and trains in our transportation network? Here we use the term trains to represent transportation running on a fixed schedule, as opposed to buses which arrive according to a random distribution.¹ We demonstrate how to solve this problem in the case where time is discretized and interarrival times for buses are given by discrete geometric random variables. Note that discrete geometric random variables

¹We acknowledge that our use of the terms buses and trains may be inaccurate for practice; still, they are useful.

provide natural approximations for continuous exponential random variables, where the accuracy of the approximation depends on the granularity of the time interval for the discretization. Hence our result can be used to approximate the continuous Poisson arrival case. Although our solution is polynomial in the number of time steps modeled, we believe it may be effective for problems of a reasonable size.

1.1 Related work

The earliest reference we have found to bus network problems is by Hall [4]. The starting point of our work is the recent paper by Datar and Ranade, who solved the problem of bus transportation networks when all bus arrivals are independent and Poisson [3]. An interesting aspect of this work is the on-line decision making process of the traveler, who chooses whether or not to take a bus as it arrives. Previous approaches required schedules that force the rider to commit to a single transit choice upon arriving at a stop, rather than flexibly choosing based on what bus gets there first [4, 9].

We also view this work as an interesting connection between algorithmic analysis and Markov decision processes. For more background on Markov decision processes, see for example [1, 6]; we offer a brief description here. In a Markov decision process, there is an underlying Markov process with associated actions and rewards. By choosing an action at a state, one affects the progress of the Markov process: the goal is to choose options that optimize the cumulative reward. In the case of the bus network problem, the actions in each state are whether or not to take a bus when it arrives, and the function we wish to optimize is the expected time to the destination. In the cases we consider here, the state also explicitly includes a time component, and hence it fits into the framework of time-dependent Markov decision processes introduced in [2]. Our work demonstrates that under certain probabilistic assumptions, there are efficient algorithms to determine the actions that yield an optimal solution in the bus network setting. Our algorithms all rely on dynamic programming, which is the fundamental technique for solving problems based on Markov decision processes [1, 6].

2 Buses with IFR Waiting Times

2.1 Probability Preliminaries

For completeness we cover basic definitions and properties of distributions we will use throughout the paper. Further information can be found in texts such as [7] or [8].

We will generally assume throughout that our random variables are non-negative with absolutely continuous cumulative distribution functions² and finite means, although our results can be modified to handle other cases, including for example discrete distributions.

²For our purposes, absolutely continuous means that the first derivative exists almost everywhere.

For a nonnegative random variable X with cumulative distribution function $F(t)$, we define the survival function to be $\bar{F}(t) = 1 - F(t)$. Formally, X is said to have *increasing failure rate* (or be IFR) if $\log \bar{F}(t)$ is concave on the support of \bar{F} . That is, $\bar{F}(t)$ is *logconcave*. Alternatively, if $f(t) = F'(t)$ is the corresponding density function, the failure rate is $r(t) = f(t)/\bar{F}(t)$. The condition that $\log \bar{F}(t)$ is concave is equivalent to the condition that r is increasing.³ The function $r(t)$ satisfies

$$r(t) = \frac{f(t)}{\bar{F}(t)} = \lim_{\Delta t \rightarrow 0} \frac{\Pr(t < X \leq t + \Delta t \mid X > t)}{\Delta t}.$$

Informally, if X represents a time spent waiting for a bus and X has increasing failure rate, it means the probability of the bus suddenly appearing increases the longer we wait.

Similarly, X has *decreasing failure rate* (or is DFR) if $\log \bar{F}$ is convex on its support, or equivalently, X is DFR if $r(t)$ is decreasing.

The mean residual life of X at time t is defined as

$$m_X(t) = E[X - t \mid X > t].$$

For example, if X represents the time until a bus arrives, the mean residual life $m_X(t)$ represents the average time until the bus arrives, given that it has not arrived during the first t units of time. Note that $m_X(t)$ is defined to be 0 where $\bar{F}(t) = 0$. The random variable X is said to have *decreasing mean residual life* or be DMRL if $m_X(t)$ is decreasing. An interesting lemma left to the reader is that if X is IFR then it is DMRL, but the reverse need not hold.

The exponential distribution is both IFR and DFR. Uniform distributions are clearly IFR. Normal distributions can be shown to be IFR [5], as can gamma distributions with certain parameters [7]. In particular, any gamma random variable that is the sum of a finite number of exponential random variables is IFR.

2.2 Form and Computation of the Optimal Schedule

We begin with a theorem that shows the form of the optimal schedule when the waiting times for buses are IFR.

Let $T(s, d, h)$ denote the expected time to reach d from s using at most h bus changes. Similarly, let $T_b(s, d, h)$ be the expected time to reach d from s using at most h bus changes, given that the rider gets on bus b now.

We will focus on a single stop s with buses B_1, B_2, \dots, B_k stopping there. (For convenience, we do not include s in the variable description of the buses, but leave it implicit.) We will also use T_i as an implicit shorthand for $T_{B_i}(s, d, h - 1)$. We let W_i be the random variable representing the waiting time for bus i , and let $W_i(t)$ be the random variable corresponding to the remaining waiting time, $[W_i - t \mid W_i > t]$.

³Here we follow the perhaps unfortunate but apparently standard practice and use “increasing” to mean “non-decreasing” and “decreasing” to mean “non-increasing” throughout. So IFR really means the failure rate is non-decreasing, even though IFR is the standard term.

There are a few additional concerns we mention here. If a bus travels through multiple stops, we must assume that the arrival distributions of buses at each stop and the travel times from stop to stop are independent. With this framework, we may assume without loss of generality that each bus travels only to a single next stop; our results below can be modified so the rider chooses the best of several possible stops along the route if there are several stops. We will make this assumption in the theorem below. Second, suppose a bus B_i visits the stop s but the rider chooses not to take it. It is not clear what arrival distribution we should use for the next visit by a bus B_i . The distribution W_i represents the waiting time from our arrival; it is not clear that we should use the same distribution after B_i itself arrives. Theorem 1 actually holds under any distribution for the waiting time of a “re-visit” by a bus B_i .

Theorem 1 *Suppose that at every bus stop, the waiting times for the buses are independent random variables with increasing failure rate. Let B_1, B_2, \dots, B_k be the buses passing through a stop s , sorted in order of increasing T_i (expected total remaining travel time to the destination d using at most $h - 1$ further bus changes). Then the optimal travel plan from s to d using at most h bus changes has the following form: take B_1 whenever it arrives; take B_2 if it arrives before time t_2^* ; take B_3 if it arrives before time t_3^* ; and so on, where the t_i^* are decreasing ($\infty \geq t_2^* \geq t_3^* \geq \dots \geq 0$).*

Proof:

We first provide the important intuition. It is clear that in the optimal schedule, bus B_1 is taken whenever it arrives, since the expected time to reach d by taking any other bus must be at least as great as T_1 .

When bus B_2 arrives, however, the best plan may involve trying to wait for bus B_1 . Clearly, the rider should wait for bus B_1 if the expected time to wait for and then take B_1 to reach d is less than the expected time if the rider now takes B_2 . That is, suppose B_2 arrives at time t , and

$$T_2 > T_1 + E[W_1(t)]; \quad (1)$$

then it is better to wait for bus B_1 . (Note that we have used in equation (1) that the waiting time for bus B_1 is independent of the arrival of bus B_2 .) The reverse is less clear; even if $T_2 < T_1 + E[W_1(t)]$, perhaps it could be better on average to wait for a following bus, hoping that it is B_1 but settling for B_3 or B_4 if we are unlucky. In fact this is not the case; we will show that the condition

$$T_2 \leq T_1 + E[W_1(t)] \quad (2)$$

is sufficient as well as necessary for taking bus 2 at time t . Using this equivalence, and the fact that $E[W_1(t)]$ is decreasing in t (since W_1 is assumed IFR), we can conclude that there is a threshold time t_2^* such that the rider should take bus B_2 if it arrives before t_2^* , where

$$t_2^* = \inf\{t : T_2 \leq T_1 + E[W_1(t)]\}.$$

Note that at times where there is equality in the above expression, either waiting or taking the bus yields the same expected time, and hence without loss of generality we may say that the optimal schedule takes B_2 if and only if it arrives before t_2^* . The argument for other buses will be similar, using induction on the B_i .

To show that condition (2) is sufficient seems difficult, since ostensibly we need to consider all possible other plans and arrival patterns of buses. We avoid this complexity by introducing an *option argument*. Let us suppose that when bus 2 arrives, we give the rider an option to force bus B_2 to wait; the rider can then board B_2 and have it leave at his or her discretion, or board another bus that arrives later. It is clear that this added option only helps the rider. Moreover, for any plan in the original setting where the rider waits for some other bus B_i with $i \geq 2$ and boards that bus, there is a plan at least as good in the option setting where the rider exercises the option and takes bus B_2 at the time it would have taken the other bus. Hence we need only consider whether the rider should take B_2 now, exercise the option (taking B_2 in the future), or wait for bus B_1 .

In this context, however, choosing to take bus B_2 in the future can never be optimal. This follows again from the fact that $E[W_1(t)]$ is decreasing, so the longer B_2 sits idle, the more appealing B_1 becomes. Therefore, the only two potentially optimal choices are to board and take B_2 immediately, or to commit to waiting for B_1 . This decision is precisely the test of Equation 2, resulting in the simple outcome that B_2 should be taken if and only if it arrives before time t_2^* .

Now let us consider the similar inductive argument for B_j , where $j > 2$. Let Z_m be the random variable representing the time to reach d using at most h bus changes, if the rider waits for one of buses B_1, B_2, \dots, B_m and uses the optimal policy for these m buses. Similarly, let $Z_m(t)$ be the time to reach d after having already waited t seconds at s . We know the form of the optimal policy on $j - 1$ buses via the inductive hypothesis. Clearly it is necessary that

$$T_j \leq E[Z_{j-1}(t)] \tag{3}$$

for it to be optimal for the rider to take bus j if it arrives at time t . To show that (3) is also sufficient, it suffices to show that Z_{j-1} is DMRL from the option argument.

We use the fact that the distribution of the W_i are IFR to show that Z_{j-1} is DMRL. Unfortunately, a direct argument is somewhat difficult, as a natural expression for $E[Z_{j-1}(t)]$ is difficult to write: the buses involved with the calculation of $Z_{j-1}(t)$ change with t . (The correct expression is therefore a sum, split according to the condition of when the first relevant bus arrives.)

We instead show that Z_{j-1} is DMRL over successive intervals. Inductively, it suffices to consider the interval $[0, t_{j-1}^*]$. The argument is simplified by constructing a new random variable Y_{j-1} , which is similar to Z_{j-1} except for the following changes. First, we replace the waiting time distribution for bus $j - 1$ by a distribution that is equal to W_{j-1} for all $t \leq t_{j-1}^*$ and is t_{j-1}^* with all remaining probability. That is, for the variable Y_{j-1} we assume that bus $j - 1$ arrives at time t_{j-1}^* if it has not otherwise arrived. Note that $E[Y_{j-1}(t)] = E[Z_{j-1}(t)]$

over the interval $[0, t_{j-1}^*]$, as this change does not affect the expected travel time over this interval. Second, for Y_{j-1} we assume that if the bus B_i is boarded, the remaining travel time is exactly the expectation T_i instead of a random variable. Again, with this change we still have $E[Y_{j-1}(t)] = E[Z_{j-1}(t)]$ over the interval $[0, t_{j-1}^*]$ (by linearity of expectations). Hence it suffices to show that Y_{j-1} is IFR to prove Z_{j-1} is DMRL.

Note that

$$\Pr(Y_{j-1} \geq x) = \Pr(W_{j-1} \geq x - T_{j-1}) \cdot \Pr(W_{j-2} \geq x - T_{j-2}) \cdot \dots \cdot \Pr(W_1 \geq x - T_1).$$

But the survival functions of every term in the product on the right hand side are logconcave in x , since the W_i are IFR. Hence the left hand side is logconcave in x , and since the left hand side is the survival function of Y_{j-1} , we have that Y_{j-1} is IFR. Hence inductively Z_{j-1} is DMRL and the optimal policy has the form given in the statement of the theorem.

Finally, note that $t_i^* \leq t_{i-1}^*$ since the $Z_i(t)$ are decreasing in i and the T_i are increasing in i . \square

Theorem 1 immediately provides an “elementary” proof of the main result by Datar and Ranade (Lemma 3.1 of [3]).

Corollary 1 *When arrivals for all buses are Poisson, then the optimal schedule has the following form: take one of buses B_1, B_2, \dots, B_j as soon as it arrives.*

Proof: In this case, the $W_i(t)$ are independent of t , so in the proof of Theorem 1 we must have that the t_i^* are all infinity or 0. \square

In the case of Poisson arrivals, Datar and Ranade show that the optimal schedule and the resulting expected travel times can be computed exactly efficiently [3]. Theorem 1 also suggests a natural way of computing an optimal schedule for our more general setting. Let Q be the maximum number of buses that pass through a station and S be the number of stations. We may compute optimal plans involving at most h bus changes inductively. This first involves sorting the buses at each station according to the time to reach the destination using $h - 1$ further bus changes. Then we compute successive values of t_i^* for each stop.

For distributions more complex than the exponential, computing the t_i^* is non-trivial. It requires computing the expected time to reach the destination using buses B_1, B_2, \dots, B_{i-1} , which may require multiple integrations over the corresponding distributions (to find the distribution of the time the first of these buses that the rider will take arrives and the corresponding probability for each bus; note the time the bus arrives and which bus it is are correlated in our case!). In practice we expect computing the t_i^* would be done numerically to suitably high precision, or possibly even by Monte Carlo simulation. We believe that the numerical analysis issues are outside the scope of this paper. Hence we simply assume the existence of a “black box” calculator for computing the t_i^* . Given this, we have the following corollary:

Corollary 2 *The optimal travel plan under the conditions of Theorem 1 can be computed in polynomial time, assuming a black box for computing the t_i^* .*

Proof: The total work from the sorting is $O(hSQ \log Q)$, and there will be $O(hSQ)$ computations using the black box. \square

2.3 Are IFR distributions necessary?

One might ask if IFR distributions are required for the result of Theorem 1. The proof itself requires only that the time to reach the destination if the rider waits for lower numbered buses is DMRL; other sets of distributions may have this property without being IFR. One natural suggestion is that perhaps if the W_i are simply DMRL then this is sufficient. We show, however, that Theorem 1 does not hold under this condition.

Theorem 2 *The results of Theorem 1 fail to hold if the W_i are only DMRL.*

Proof: We construct a counterexample. Let X be uniform over the range $[0, 2] \cup [4, 12]$. It is simple to check that X is DMRL. Now suppose there are two types of buses traveling from point a to point b , each with waiting distribution X . Fast buses have a constant travel time of 1; slow buses have a constant travel time of 2. Clearly if a fast bus come the rider should always take it. Intuitively, however, if there are enough fast buses, the rider should not take a slow bus that arrives early, because it is likely that a fast bus will shortly come. Once the rider has waited almost two time units, however, he should take a slow bus if it comes, since no buses arrive in the interval $[2, 4]$ and otherwise he is likely to end up waiting substantially before a fast bus appears. A calculation shows that if there is one slow bus and twelve fast buses, we should take a slow bus only if it arrives in the interval $[0.558, 2]$. \square

The counterexample of Theorem 2 can be modified in various ways. For example, we can change the distribution X so that its support is a closed interval by adding a small ϵ weight over the interval $[2, 4]$. Also, by considering distributions X that consist of more disjoint intervals, we can construct examples where the proper times to take the slow bus consist of two or more disjoint intervals. The point behind Theorem 2 is that even though each fast bus has decreasing mean time to live, the random variable for the time until the first fast bus arrives does not. That is, the family of random variables with decreasing mean time to live is not closed under minimization, while IFR random variables are.

An interesting open question this counterexample raises is whether there is a natural way to relate the complexity of the waiting time distributions and the complexity of the form of the optimal schedule.

3 Buses and Trains

We now consider another issue suggested in the conclusion of [3], and also examined in [2]: networks with mixed forms of transportation, such as buses and trains. Recall that in our model buses have an associated random waiting time distribution and an associated

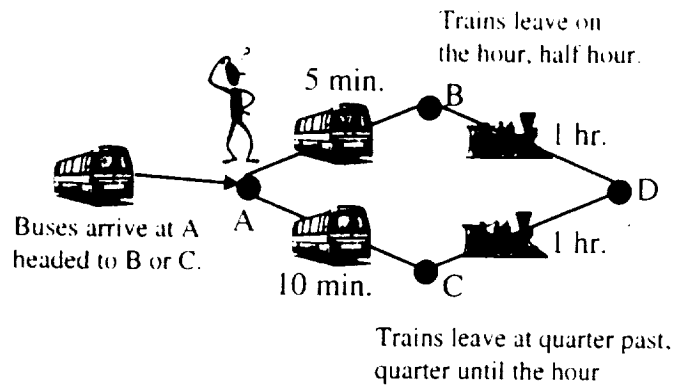


Figure 1: A basic bus and train network. Arrivals of buses from A to B and from A to C are each Poisson with an average wait of ten minutes. Travel times are constant. Trains from B leave on the hour and the half hour; trains from C leave fifteen and forty-five minutes into the hour.

random travel time distribution. We shall use the term *train* to refer (metaphorically) to transportation that arrives and departs at fixed absolute times. For example, consider Figure 1. From station A, the rider may catch a bus to either station B or station C. We assume the travel time from A to B is a constant five minutes and the travel time from A to C is a constant ten minutes. Arrivals of buses that travel from A to B are a Poisson process, with an average waiting time of ten minutes; the same holds for buses from A to C. At both stations B and C there are trains that run to station D, with the travel time on the train being one hour. Trains from B leave on the hour and the half hour; trains from C leave fifteen and forty-five minutes into the hour.

This simple example highlights that introducing trains leads to substantial difficulties. Of primary importance is the introduction of absolute time; we are not only concerned with how long the rider has spent at the bus station, as in the problem with only buses, but the actual time until the trains depart. The expected time to reach the destination D from train stations B and C is not constant, as it was in the pure-bus setting of Section 2, but depends on the time the rider arrives at D. Because of this, the ideas behind Theorem 1 no longer directly apply. In particular, there are times where the rider should pass up the bus to B in order to wait for the bus to C, and other times when the rider should do the opposite. For example, if a bus to station B arrives at station A just four minutes before the hour, we know that taking the bus will cause us to wait at station B. The rider is better off waiting for a bus to station C, and possibly catching a bus to station B later if necessary.

In this section, we present an approach for handling mixed networks of buses and trains in the case where bus arrival times are discrete geometric random variables, which can be used to approximate the case of the continuous Poisson arrival process. Our method will

be to set up the problem as a large dynamic programming problem, or equivalently, as a Markov decision process. (Dynamic programming is a standard technique for Markov decision processes; again, see [1, 6].) We first present our approach via the example above, and then discuss the general framework for larger problems.

We first clarify here why we limit the bus arrival processes to be discrete geometric random variables. From our example, we can see that it is possible in mixed bus and train networks that the rider chooses not to take a bus at some time, only to take a bus on the same route later. If the bus arrival process is a geometric random variable, the fact that a bus has previously arrived need not be recorded in the state space: we may forget that the bus has arrived, as it does not affect the arrival of future buses. (This is the memorylessness property of geometric and exponential random variables.) If, however, a bus has a more complicated arrival process, then the last time a bus on that route arrived may be relevant information for determining the arrival of the next bus on that route. Keeping track of such information as the last arrival of each bus would lead to a more complex, higher-dimensional state. Although handling such a state is theoretically feasible using the techniques we suggest, we do not address this issue here.

For our problem, the state space will be pairs (s, t) , where s is a station and t is the current absolute time. To be at the state (s, t) denotes that the rider is still waiting at station s at time t . In order to make the underlying state space countable, we must assume time is discretized. Moreover, for the state space to be finite, we must also assume an ending time for the process. For example, we may assume that the buses and trains start running at noon and stop running at midnight, at which point one must call a friend for a ride. To penalize this action, we make the cost associated with it very high but finite (such as two hours).

Discrete geometric random variables can naturally be used to approximate continuous exponentially distributed random variables: the error in the approximation depends on the granularity of the discrete time scale. Hence this approach can be used to approximate behavior when bus arrival processes are Poisson. On the other hand, the number of states required is proportional to the number of discrete time steps being modeled.

Finally, for convenience we will assume here that each bus travels from our current stop to a unique other stop as opposed to multiple stops. The case where buses have multiple stops can be handled in a similar fashion (with a possible increase in the size of the state space).

3.1 The Dynamic Program

We first consider our example. For convenience let us assume that time is discretized in minutes, and buses leave on the minute. Hence for example the rider may begin at state (A,11:59 am), and if a bus arrives in the intervening minute, he may get on the bus and leave station A at 12:00. Of course the rider may choose not to get on the bus, in which case the rider will be at state (A,12:00 am). We wish to optimize the expected time of arrival at station D.

The possible actions at each state consist of the *list* of buses we will take if such a bus arrives at the station over the next minute from that time. We assume in this discretized version that buses may arrive in the same interval, so that our possible action at each state is a sorted list of buses that we will take if a bus arrives, with the sorted order giving a preference if two buses arrive at the same time.

In our simple example, let $E(s, t)$ be the expected time to reach D from state s at time t . Note that $E(B, t)$ and $E(C, t)$ are trivial to compute. Let $z = 1 - \exp(-0.1)$ be the probability that a bus headed for B (or equivalently for C) arrives at A during a minute. We have the recurrence

$$\begin{aligned} E(A, t) = & (1 - z)^2 E(A, t + 1) + \\ & z(1 - z) \min(E(A, t + 1), E(B, t + 6)) + \\ & z(1 - z) \min(E(A, t + 1), E(C, t + 11)) + \\ & z^2 \min(E(A, t + 1), E(B, t + 6), E(C, t + 11)). \end{aligned}$$

We solve this recurrence for decreasing t . From the recurrence we can naturally derive the correct actions: for example, if $E(A, t + 1) < E(B, t + 6)$ then we will not take the bus to B. The results for our example are given in Table 1.

Time	0-2	3-4	5-21	22-24	25-30
Action	C	C,B	B	B,C	C

Table 1: At more than a few hours from the end of the day, the optimal strategy has a half hour cycle. At states where t is 0 to 2 minutes over the half hour, plan only to take bus C if it arrives over the next minute.

Theorem 3 *The travel plan for optimizing the expected travel time for networks with buses and trains, where buses have discrete geometric arrival distributions, can be computed in time polynomial in the number of stops, the maximum number of buses and trains at a stop, and the total number of time units simulated. This holds even if bus travel times are random variables that depend on the time of arrival to the station.*

Proof: We provide a more general framework and corresponding bounds on the time to compute the optimal strategy. We use dynamic programming, computing the $E(s, t)$ in reverse temporal order; that is, we start at the end of the process, and compute $E(s, t)$ for all s (in any order) using already computed values $E(s, u)$ with $u > t$. Suppose that a maximum of Q buses or trains pass through any of S total stops, and our process lasts for T units of time. For each time state (s, t) , there are at most $\sum_{i=0}^Q \binom{Q}{i} i!$ possible actions, as each ordered subset of the buses and trains are a possible action. However, following the idea of Theorem 1 we can simplify considerably by sorting the buses and trains by the expected time to reach the destination if we choose that option at that time. Every action that is better than waiting at the current stop (i.e, better than $E(s, t + 1)$) is one that will

be taken, and the sorted order provides the preference. Note that we can sort the transport options, but the results may be different for different time steps.

In our example, we have that the bus travel times are constant. In this case, once we have the sorted order, computing $E(s, t)$ can be done in time $O(Q)$ by considering the arrival possibilities in sorted order. For convenience suppose there are only buses at the station (trains are easy to handle, as they are either ready to leave or not). If the first bus arrives, we take it; if not, but the second bus arrives, we take that; and so on. There are only $O(Q)$ possibilities to consider. Hence the total time to compute optimal schedules in this case is $O(STQ \log Q)$.

If instead bus travel times are given by a fixed discrete random distribution, or even a discrete distribution that varies over time, this only increases the work to compute the expected times to reach the destination by a factor of $O(T)$, for total work $O(STQ \log Q + SQT^2)$.

(Note: in the case where bus travel times are given by fixed discrete random variables, standard convolution techniques may reduce the total work to $O(STQ \log Q + SQT \log T)$; however, it appears some additional assumptions are necessary for these methods to apply. We will explain further in the final version of the paper.) \square

Although the complexity of these solutions may be large when computing over long time intervals, they appear feasible for reasonable-sized systems. We also note that another advantage of this setup is that we can handle value functions more general than the expected travel time: for example, we could use the same approach to maximize the probability of reaching our destination by a certain time.

To summarize, this framework improves over previous work in the following respects. In comparison to the work of [3], we show that handling buses with Poisson arrivals and trains is possible; moreover, we show that the simple form of the optimal schedule we have shown in Section 2 is not possible in this setting. In comparison with previous work on Markov decision processes such as [2, 9], we have shown how to handle the problem of waiting for multiple buses at a station in the case of a geometric arrival process, which leads to a relatively simple state space.

4 Conclusions and Future Work

We have expanded previous work on stochastic transportation networks in two ways. First, we provided an algorithm for finding optimal schedules for bus networks where bus arrival distributions have increasing failure rate. Second, we have given an algorithm for finding optimal schedules in mixed networks of buses and trains when the bus arrival distributions are discrete geometric random variables. We plan to implement these algorithms and test them on artificially generated and real data in the near future.

There remain many open questions to pursue; we suggest two here. First, fast approximation algorithms would be useful, especially for transportation networks that change often. Moreover, approximation algorithms may allow more general distribution classes to

be handled. Second, we might consider the situation when the transportation network may provide additional information. For example, buses equipped with global positioning equipment and wireless communication may be able to provide their position. In this situation, a rider determining whether or not to get on a bus may have more detailed information available about the waiting time for other buses.

References

- [1] D. Bertsekas and J. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA, 1996.
- [2] J. A. Boyan and M. L. Littman. Exact solutions to time-dependent MDPs. Submitted to Neural Information Processing Systems, 2000.
- [3] M. Datar and A. Ranade. Commuting with delay prone buses. In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 22–29, 2000.
- [4] R. W. Hall. The fastest path through a network with random time-dependent travel times. *Transportation Science*, 20(3):182–188, 1986.
- [5] K. C. Kapur and L. A. Lamberson. *Reliability in Engineering Design*. John Wiley and Sons, New York, 1977.
- [6] M. L. Puterman. *Markov Decision Processes—Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, 1994.
- [7] S. Ross. *Stochastic Processes*. John Wiley and Sons, New York, 1983.
- [8] D. Stoyan. *Comparison Methods for Queues and Other Stochastic Models*. John Wiley and Sons, Berlin, 1983.
- [9] M. P. Wellman, M. Ford, and K. Larson. Path planning under time-dependent uncertainty. In *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence (UAI'95)*, pages 532–539. Morgan Kaufmann, 1995.